

1662-36900
P00-3156

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR UNITED STATES LETTERS PATENT

**SIMULTANEOUS AND REDUNDANTLY THREADED
PROCESSOR STORE INSTRUCTION COMPARATOR**

By:

Shubhendu S. Mukherjee
43 Gates Street
Framingham, Massachusetts 01702
Citizenship: India

Steven K. Reinhardt
3158 Kilburn Park Cir.
Ann Arbor, Michigan 48105
Citizenship: U.S.A.

**SIMULTANEOUS AND REDUNDANTLY THREADED
PROCESSOR STORE INSTRUCTION COMPARATOR**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a non-provisional application claiming priority to provisional application Serial No. 60/198,530, filed on April 19, 2000, entitled "Transient Fault Detection Via Simultaneous Multithreading," the teachings of which are incorporated by reference herein as if reproduced in full below.

[0002] This application also relates to application Serial No. 09/584,034, filed May 30, 2000, entitled "Slack Fetch to Improve Performance in a Simultaneous and Redundantly Threaded Processor," the teachings of which are incorporated by reference herein as if reproduced in full below.

[0003] This application also relates to application Serial No. _____, entitled "Simultaneous and Redundantly Threaded Processor Uncache Load Address Comparator and Data Value Replication Circuit," (Attorney Docket No. 1662-37400) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

[0004] This application also relates to application Serial No. _____, entitled "Cycle Count Replication in a Simultaneous and Redundantly Threaded Processor," (Attorney Docket No. 1662-37000) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

[0005] This application also relates to application Serial No. _____, entitled "Active Load Address Buffer," (Attorney Docket No. 1662-37100) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

[0006] This application also relates to application Serial No. _____, entitled "Simultaneous and Redundantly Threaded Processor Branch Outcome Queue," (Attorney Docket No. 1662-37200) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

[0007] This application also relates to application Serial No. _____, entitled "Input Replicator for Interrupts in a Simultaneous and Redundantly Threaded Processor," (Attorney Docket No. 1662-37300) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

[0008] This application also relates to application Serial No. _____, entitled "Load Value Queue Input Replication in a Simultaneous and Redundantly Threaded Processor," (Attorney Docket No. 1662-37500) filed concurrently herewith, the teachings of which are incorporated by reference herein as if reproduced in full below.

**STATEMENT REGARDING FEDERALLY SPONSORED
RESEARCH OR DEVELOPMENT**

[0009] Not applicable.

BACKGROUND OF THE INVENTION

Field of the Invention

[0010] The present invention relates generally to microprocessors. More particularly, the present invention relates to a pipelined, simultaneous and redundantly threaded processor adapted to execute the same instruction set in at least two separate threads, and to verify for fault detection purposes only memory requests that affect data values in system memory. More particularly still,

the invention relates to detecting transient faults between multiple processor threads by comparison of their store instructions.

Background of the Invention

[0011] Solid state electronics, such as microprocessors, are susceptible to transient hardware faults. For example, cosmic radiation can alter the voltage levels that represent data values in microprocessors, which typically include tens or hundreds of thousands of transistors. The changed voltage levels change the state of individual transistors, causing faulty operation. Faults caused by cosmic radiation typically are temporary and the transistors eventually operate normally again. The frequency of such transient faults is relatively low -- typically less than one fault per year per thousand computers. Because of this relatively low failure rate, making computers fault tolerant currently is attractive more for mission-critical applications, such as online transaction processing and the space program, than computers used by average consumers. However, future microprocessors will be more prone to transient fault due to their smaller anticipated size, reduced voltage levels, higher transistor count, and reduced noise margins. Accordingly, even low-end personal computers benefit from being able to protect against such faults.

[0012] One way to protect solid state electronics from faults resulting from cosmic radiation is to surround the potentially effected electronics by a sufficient amount of concrete. It has been calculated that the energy flux of the cosmic radiation can be reduced to acceptable levels with at least six feet concrete surrounding the chips to be protected. For obvious reasons, protecting electronics from faults caused by cosmic radiation with six feet of concrete usually is not feasible as computers are usually placed in buildings that have already been constructed without this amount of concrete. Because of the relatively low occurrence rate, other techniques for protecting

microprocessors from faults created by cosmic radiation have been suggested or implemented that merely check for and correct the transient failures when they occur.

[0013] Rather than attempting to create an impenetrable barrier through which cosmic rays cannot pierce, it is generally more economically feasible and otherwise more desirable to provide the effected electronics with a way to detect and recover from faults caused by cosmic radiation. In this manner, a cosmic ray may still impact the device and cause a fault, but the device or system in which the device resides can detect and recover from the fault. This disclosure focuses on enabling microprocessors (referred to throughout this disclosure simply as "processors") to recover from a fault condition.

[0014] One technique for detecting transient faults is implemented in the Compaq Himalaya system. This technique includes two identical "lockstepped" microprocessors that have their clock cycles synchronized, and both processors are provided with identical inputs (*i.e.*, the same instructions to execute, the same data, etc.). In the Compaq Himalaya system, each input to the processors, and each output from the processors, is verified and checked for any indication of a transient fault. That is, the hardware of the Himalaya system verifies all signals going to and leaving the Himalaya processors at the hardware signal level -- the voltage levels on each conductor of each bus are compared. The hardware performing these checks and verifications is not concerned with the particular type of instruction it is comparing; rather, it is only concerned that two digital signals match. Thus, there is significant hardware and spatial overhead associated with performing transient fault detection by lockstepping duplicate processors in this manner.

[0015] The latest generation of high-speed processors achieve some of their processing speed advantage through the use of a "pipeline." A "pipelined" processor includes a series of units (*e.g.*, fetch unit, decode unit, execution units, etc.), arranged so that several units can simultaneously

process an appropriate part of several instructions. Thus, while one instruction is decoded, an earlier fetched instruction is executed. These instructions may come from one or more threads. Thus, a “simultaneous multithreaded” (“SMT”) processor permits instructions from two or more different program threads (e.g., applications) to be processed simultaneously. However, it is possible to cycle lockstep the threads of an SMT processor to achieve fault tolerance.

[0016] An SMT processor can be modified so that the same program is simultaneously executed in two separate threads to provide fault tolerance within a single processor. Such a processor is called a simultaneous and redundantly threaded (“SRT”) processor. Some of the modifications to turn a lockstep SMT processor into an SRT processor are described in Provisional Application Serial No. 60/198,530. However, to utilize known transient fault detection requires that each thread of the SRT processor be lockstepped (as opposed to having two SRT processors lockstepped to each other). Hardware within the processor itself (in the Himalaya, the hardware is external to each processor) verifies the digital signal on each conductor of each bus. While increasing processor performance and yet still doing transient fault protection in this manner may have advantages over previous fault detecting systems, SRT processor performance can be enhanced.

[0017] One such performance enhancing technique is to allow each processor to run independently. More particularly, one thread is allowed to execute program instructions ahead of the second thread. In this way, memory fetches and branch speculations are resolved ahead of time for the trailing thread. However, verifying, at the signal level, each input and output of each thread becomes complicated when the threads are not lockstepped (executing the same instruction at the same time).

[0018] A second performance enhancing technique for pipelined computers is an “out-of-order” processor. In an out-of-order processor each thread need not execute the program in the order it is

presented; but rather, each thread may execute program steps out of sequence. The technique of fault tolerance by verifying bus voltage patterns between the two threads becomes increasingly difficult when each thread is capable of out-of-order processing. The problem is further exacerbated if the one processor thread leads in overall processing location within the executed program. In this situation not only would the leading thread be ahead, but this thread may also execute the instructions encountered in a different sequence than the trailing thread.

[0019] The final performance enhancing technique of SRT processor is speculative branch execution. In speculative branch execution a processor effectively guesses the outcome of a branch in the program thread and executes subsequent steps based on that speculation. If the speculation was correct, the processor saves significant time (over stalling for example) until the branch decision is resolved. In the case of an SRT processor it is possible that each thread makes speculative branch execution different than the other. Thus, it is impossible to do transient fault protection using known techniques – verifying digital signals on each bus – because it is possible there may be no corresponding signal between two threads.

[0020] What is needed is an SRT processor that can achieve performance gains over an SRT processor in which each thread is lockstepped by using the performance enhancing techniques noted above, and that can also do transient fault detection.

BRIEF SUMMARY OF THE INVENTION

[0021] The problems noted above are solved in large part by a simultaneous and redundantly threaded processor that has performance gains over an SRT processor with lockstepped threads and provides transient fault tolerance. The processor checks for transient faults by checking only memory requests (input/output (“I/O”) commands, I/O requests) that directly or indirectly affect

data values in system memory. More particularly, the preferred embodiments verify only writes (stores) that change data outside the bounds of the processor and uncached reads, e.g., a read from a virtual address space mapped to the hardware of an I/O device. Because this transient fault detection does not need to verify every input and output at the signal level, the transient fault protection extends to the threaded “out-of-order” processors, processors whose threads perform independent speculative branch execution, and processors with leading and lagging thread execution.

[0022] A preferred embodiment of the invention comprises a store queue and a compare circuit. The processor thread executing the program ahead, the leading thread, writes its committed stores to the store queue. Subsequently, the processor thread lagging or trailing, the trailing thread, writes its committed stores to the queue. Because the processors are preferably removed from each other in program execution stage and also may be executing the programs in different orders in each thread, there may be many committed stores written to the queue between when the leading queue writes and when the trailing queue writes the corresponding store. A compare circuit periodically scans the store queue looking for the corresponding stores. If the address and data information of the corresponding stores matches exactly, then each of the processors have operated without fault, and the data is written to the data cache or main memory. However, if any differences exist in the address or the data (actual data or differences in size of the store) to be written, the compare circuit initiates a fault recovery sequence.

[0023] Alternatively, a second embodiment of the invention comprises the store queue into which the leading thread places its committed stores. As the trailing thread reaches this point in the program execution, hardware and firmware associated with that thread compares the committed store, without the need of placing that store in the same queue as the previous committed store, and

finds the corresponding committed store from the leading thread. If these two committed stores match exactly, the committed store placed in the queue is marked as verified and the trailing thread store is effectively discarded. The verified committed store is then sent to its appropriate location in the cache or main memory areas.

BRIEF DESCRIPTION OF THE DRAWINGS

[0024] For a detailed description of the preferred embodiments of the invention, reference will now be made to the accompanying drawings in which:

[0025] Figure 1 is a diagram of a computer system constructed in accordance with the preferred embodiment of the invention and including a simultaneous and redundantly threaded processor; and

[0026] Figure 2 is a block diagram of the simultaneous and redundantly threaded processor from Figure 1 in accordance with the preferred embodiment that includes a compare circuit to check for transient faults manifested in differences in committed stores.

NOTATION AND NOMENCLATURE

[0027] Certain terms are used throughout the following description and claims to refer to particular system components. As one skilled in the art will appreciate, microprocessor companies may refer to a component by different names. This document does not intend to distinguish between components that differ in name but not function. In the following discussion and in the claims, the terms "including" and "comprising" are used in an open-ended fashion, and thus should be interpreted to mean "including, but not limited to...". Also, the term "couple" or "couples" is intended to mean either an indirect or direct electrical connection. Thus, if a first device couples to

a second device, that connection may be through a direct electrical connection, or through an indirect electrical connection via other devices and connections.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0028] Figure 1 shows a computer system 90 including a pipelined, simultaneous and redundantly threaded (“SRT”) processor 100 constructed in accordance with a preferred embodiment of the invention. Besides processor 100, computer system 90 also preferably includes system main memory in the form of a dynamic random access memory (“DRAM”) 92, an input/output (“I/O”) controller 93, and various I/O devices which may include a floppy drive 94, a hard drive 95, a keyboard 96, and the like. Each of the various I/O devices may have on-board memory, and the combination of this I/O device memory and the system main memory makes up the overall system memory. The I/O controller 93 provides an interface between processor 100 and the various I/O devices 94-96. The DRAM 92 can be any suitable type of memory devices such as RAMBUS™ memory. In addition, SRT processor 100 may also be coupled to other SRT processors if desired in a commonly known “Manhattan” grid, or other suitable architecture.

[0029] Figure 2 shows the SRT processor 100 of Figure 1 in greater detail. Referring to Figure 2, processor 100 preferably comprises a pipelined architecture which includes a series of functional units, arranged so that several units can be simultaneously processing appropriate parts of several instructions. Fetch unit 102 uses a program counter 106 for assistance as to which instruction to fetch. Being a multithreaded processor, the fetch unit 102 preferably can simultaneously fetch instructions for multiple thread execution. A separate program counter 106 is associated with each thread. Each program counter 106 is a register that contains the address of the next instruction to be fetched from the corresponding thread by the fetch unit 102. Figure 2 shows two program

counters 106 to permit the simultaneous fetching of instructions from two threads. It should be recognized, however, that additional program counters can be provided to fetch instructions from more than two threads simultaneously.

[0030] Fetch unit 102 includes branch prediction logic 103 which permits the fetch unit 102 to speculate ahead on branch instructions. In order to keep the pipeline full (which is desirable for efficient operation), the branch predictor logic 103 speculates the outcome of a branch instruction before the branch instruction is actually executed. Branch predictor 103 generally bases its speculation on previous instructions. Any suitable speculation algorithm can be used in branch predictor 103. Also, each thread preferably has its own branch prediction unit 103 (not shown).

Referring still to Figure 2, instruction cache 110 preferably provides a temporary storage buffer for the instructions to be executed. Decode logic 114 preferably retrieves the instructions from instruction cache 110 and determines the type of each instruction (*e.g.*, add, subtract, load, store, etc.). Decoded instructions are then preferably passed to the register rename logic 118, which maps logical registers onto a pool of physical registers.

[0031] The register update unit (“RUU”) 130 provides an instruction queue for the instructions to be executed. The RUU 130 serves as a combination of global reservation station pool, rename register file, and reorder buffer. The RUU 130 breaks load and store instructions into an address portion and a memory (*i.e.*, register) reference. The address portion is placed in the RUU 130, while the memory reference portion is placed into a load/store queue (not specifically shown in Figure 2).

[0032] The floating point register 122 and integer register 126 are used for the execution of instructions that require the use of such registers as is known by those of ordinary skill in the art. These registers 122, 126 can be loaded with data from the data cache 146. The registers also

provide their contents to the RUU 130. Figure 2 shows two sets of floating point registers 122 and integer registers 126 for a two-thread processor. However, each thread of the microprocessor preferably has its own set of floating point registers 122 and integer registers 126, thus multiple sets of these registers may be present, depending on the number of threads of the processor.

[0033] The execution units 134, 138, and 142 comprise a floating point execution unit 134, a load/store execution unit 138, and an integer execution unit 142. Each execution unit performs the operation specified by the corresponding instruction type. Accordingly, the floating point execution units 134 execute floating instructions such as multiply and divide instruction while the integer execution units 142 execute integer-based instructions. The load/store units 138 perform load operations in which data from memory is loaded into a register 122 or 126. The load/store units 138 also perform store operations in which data from registers 122, 126 is written to data cache 146 and/or DRAM memory 92 (Figure 1). Operation of the load/store units 138 of the preferred embodiments are discussed more fully below.

[0034] The architecture and components described herein are typical of microprocessors, and particularly pipelined, multithreaded processors. Numerous modifications can be made from that shown in Figure 2. For example, the locations of the RUU 130 and registers 122, 126 can be reversed if desired. For additional information, the following references, all of which are incorporated herein by reference, may be consulted for additional information if needed: U.S. Patent Application Serial No. 08/775,553, filed December 31, 1996, and "Exploiting Choice: Instruction Fetch and Issue on an Implementable Simultaneous Multithreaded Processor," by D. Tullsen, S. Eggers, J. Emer, H. Levy, J. Lo and R. Stamm, Proceedings of the 23rd Annual International Symposium on Computer Architecture, Philadelphia, PA, May 1996.

[0035] An embodiment provides a performance enhancement to SRT processor. The preferred SRT processor 100 is capable of processing instructions from two different threads simultaneously. Such a processor in fact can be made to execute the same program in the two different threads. More particularly, an SRT processor of an embodiment preferably executes the same program in each thread, however, one thread preferably leads the program execution, the leading thread, and likewise the second thread trails the program execution, the trailing thread. Performance gains may be had over a lockstepped SRT in this manner by having data reads and branch predictions resolved before the second thread reaches the program execution stages where those pieces of information are requested or need to be known. For further information on an embodiment to achieve these performance gains, see co-pending application Serial No. 09/584,034 titled "Slack Fetch To Improve Performance in Simultaneous and Redundantly Threaded Processor," filed May 30, 2000. Processing the same program through the processor in two different threads permits the processor to detect transient faults caused by cosmic radiation as noted above.

[0036] Transient fault detection is preferably accomplished by checking or verifying only memory requests that directly or indirectly affect data values in memory. More particularly, and referring to Figure 2, an embodiment comprises a store queue 140 and a compare circuit 148. The leading thread of the SRT processor preferably writes its committed stores to the store queue 140. A committed store is a write request generated during execution of a program in the thread. The "committed" portion of the committed store indicates that the store or write request is an affirmative request, as opposed to a write request that may be generated from program steps in a branch of the program that was speculatively executed. A store or write request generated by program steps from the speculative portion of a program may not come to fruition if the branch prediction was incorrect. Thus, a committed store is a write request from a portion of the program

that was non-speculatively executed, or from a speculative portion of the program that was indeed the correct branch for execution.

[0037] Thus, the leading thread preferably writes this committed store to the store queue 140. However, the committed store does not execute upon its initial insertion in the store queue 140. Rather, the committed store waits in the queue for the trailing thread to reach that point in the program execution. When the trailing thread reaches that point, this second thread then preferably writes its committed store to the store queue 140. In one embodiment, the threads then continue about their execution without regard to the remaining steps required to execute the committed store. In the situation where each of the threads places their committed store into the store queue 140, compare circuit 148 preferably comes into play.

[0038] Compare circuit 148 preferably periodically scans the content of the store queue 140. Compare circuit 148 looks for matching store requests. More particularly, the compare circuit 148 preferably compares the address and data from each related committed store request from each thread. The comparison may be a bit for bit comparison of the data, but may also include comparing the sizes of the corresponding stores. If these committed stores from each thread match exactly (their address and write data are exactly the same), then only one of those committed stores or write request is allowed to proceed to update the cache or main memory.

[0039] If, however, the compare circuit 148 determines that corresponding committed store or write requests are different in some respect, then a transient fault has occurred. That is, if the address or data of corresponding committed stores are different, then a transient fault has occurred in one of the processor threads. In this situation, the compare circuit 148 preferably initiates a fault recovery scheme or sequence. This fault recovery scheme preferably includes restarting each of the microprocessor threads at a point in the program before the fault occurred.

[0040] In a second embodiment of the invention, the trailing thread has sufficient hardware and firmware to perform the verification of the committed store itself. In this second embodiment, the leading thread writes its committed store to the store queue 140. At some time thereafter, the trailing thread commits its corresponding store and, rather than simply placing it in the store queue 140, trailing thread logic scans the store queue 140 for the corresponding write request. If the address and data of the committed store written previously by the leading thread exactly matches the committed store generated by the trailing thread, the leading thread store in the store buffer 140 is validated, for example, by setting a valid bit within the queue, and the load/store unit 138 therefore executes the data write.

[0041] Accordingly, the preferred embodiment of the invention provides for transient fault detection of a SRT processor by comparing committed stores prior to their writes to data cache or system main memory. The above discussion is meant to be illustrative of the principles and various embodiments of the present invention. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. For example, although the embodiments discussed above describe a storage queue in which at least one of the committed stores are placed prior to verification, one of ordinary skill in the art could design many ways in which these two committed stores are compared prior to their execution. That is to say, the compare circuit and storage queue are only exemplary of the idea of verifying the committed stores as a transient fault detection mechanism. Importantly, the transient fault detection for an SRT processor described above is equally applicable in processors that are capable of “out-of-order” program execution, and those that are not. The transient fault detection of the preferred embodiments is not constrained by differences in the order of execution or differences in

branch speculation. It is intended that the following claims be interpreted to embrace these and other variations and modifications.